Geospatial Systems Centre for Doctoral Training

PhD Project Proposal

Real-Time Prediction of Spatiotemporal Dynamics in the Built Environment: An Artificial Intelligence and Internet of Things Approach

Carrow Morris-Wiltshire
05/12/2023

Carrow Morris-Wiltshire is funded by the DSTL and the EPSRC Centre for Doctoral Training (CDT) in Geospatial Systems (RES/0539/7397). Special thanks are given to the supervisors of this project, Prof. Stuart Barr, Prof. Phil James, and Keith Hermiston for their support and guidance throughout the writing of this proposal.

"*The worth of idea is based on how effective it is in explaining and predicting natural phenomena*" – The Philosophy of Instrumentalism

Acronyms

IoT - Internet of Things

CAS – Complex Adaptive System

CASoS - Complex Adaptive System of Systems

DTfSA – Digital Twin for Situational Awareness

DTfSDA – Digital Twin for Strategic Decision-Making

PLM – Product Lifecycle Management

BIM – Building Information Modelling

LSTM – Long Short-Term Memory

ABM – Agent Based Model

GNN – Graph Neural-Network

ODMD – Optimised Dynamic Mode Decomposition

CCTV – Closed Circuit TeleVision

DQ – Data Quality

Contents

1 Motivation

1.1 The Complexity of Urban Problems

Navigating the complex web of challenges in modern cities, from pollution and pandemics to intricate transportation and infrastructure issues, demands a nuanced understanding of their interconnected nature (UN-Habitat, 2022). In aspiring to transform cities into inclusive, safe, resilient, and sustainable entities (United Nations, 2015), we confront the reality that cities function as complex adaptive systems (CAS), akin to living organisms (Dawson, 2011). Achieving these transformative goals requires profound comprehension of urban behaviour under planned interventions such as policy implementation or infrastructure provision, and unexpected perturbations such as extreme weather or infrastructure failure (CDBB, 2022b).

It is necessary to accurately model the city as a complex adaptive system of systems (CASoS) to achieve this (Shi *et al.*, 2021a). Here, the term ‘model’ refers to a mathematical and/or conceptual representation of a system of ideas, events, or processes used for predicting phenomena. CASoS encompass various complex adaptive systems (CAS) that exhibit adaptive, evolving, and interdependent characteristics (Baldwin *et al.*, 2011). Intervening effectively in a CASoS (city or urban region) involves first identifying a problem (such as high levels of crime) and then generating possible solutions (such as employing more police officers), using a model to estimate the solution's effectiveness. The difficulty lies in the fact that problems in a CASoS are inherently challenging to define (Rittel and Webber, 1973), making questions such as ‘why are some cities not safe?’ very difficult to answer. This complexity is compounded by the interconnected nature of CASoS, where addressing one aspect can unpredictably influence others.

The characteristics of cities (as CASoS) pose significant computational modelling challenges across the spatial and temporal dimensions (Shi *et al.*, 2021b). Events and impacts range from localised short-term effects such as the impact of flooding on daily commutes (Barr *et al.*, 2020b), to global long-term behavioural change such as the post-pandemic shift in travel demand and mode choice (Batty, 2022). Both mitigation of and adaption to these events often entail intricate trade-offs, for example, protecting cities against extreme weather events whilst minimising resource consumption and emissions. Addressing one objective can sometimes negatively impact another (Caparros-Midwood *et al.*, 2019). Many urban problems are multi-objective, multi-dimensional and spatiotemporally diverse, often described as *wicked problems* (Goodspeed, 2015).

Wicked problems, such as climate change, social injustice, and healthcare, are defined by their social complexity and lack of a determinable endpoint (Tonkinwise, 2015).

Understanding a wicked problem is difficult because it doesn't stay the same; it evolves and shifts, often in response to any attempts to understand it. Similarly, each step taken to solve a wicked problem, can change the problem. There are no definitive rules to tell when a solution to such a problem is complete, nor any ways to test it off-line (Knapp, 2015). This does not mean that we cannot try to make *positive interventions* for wicked problems (Rittel and Webber, 1973). Instead, it means that to create predictive models that are computationally realistic, the goal of intervention must shift from 'solve' to 'improve'. The wickedness of a problem is not about a higher degree of complexity, but about a fundamentally different kind of challenge to the design process, making solution secondary, and problem understanding central (Basadur *et al.*, 2007).

To comprehend the responses of urban systems to change, and thoroughly understand the problem that we are trying to improve, it becomes imperative to incorporate into our models the most unpredictable element of cities – human behaviour (Batty, 2018). State of the art approaches to urban modelling focus on the concept of a 'digital twin', a digital replica of physical assets and processes. Digital twins offer a cost-effective way to trial solutions in a virtual environment before implementing them, addressing the dynamic and interconnected nature of urban problems (Birks *et al.*, 2020; Grieves, 2022). These digital test runs of interventions in complex systems are essential in urban decision making due to the one-shot nature of wicked problems – once an intervention is made, the initial problem changes. However, in reality, most 'digital twins' are built for asset-level interventions (like building construction and maintenance), acting in isolation, and neglecting broader systemic implications arising from interactions between irrational agents (Boje *et al.*, 2020). These asset-level digital twins need to be connected to one another and integrated with models of complex city systems, including comprehensive representation and treatment of population characteristics and demography, mobility, economic activity, and environmental sustainability. This integration is vital to enable highly effective decision-making for wicked problems (CDBB, 2022b).

Assimilating current data into CASoS models is critical to predict future states with accuracy (Lewis *et al.*, 2006; Alizadeh *et al.*, 2020). Consider weather forecasting: its effectiveness is limited if based solely on a complex systems model calibrated with historical data, as such models tend to diverge from reality over time. The model's usefulness arises from its ability to assimilate the latest data and update its predictions accordingly. The process of data assimilation is illustrated in Figure 1. Data assimilation acts as a crucial bridge, continuously merging new observations with existing models to enhance the reliability and precision of future predictions. This dynamic updating ensures that the model remains closely aligned with real-world changes, maintaining its relevance and accuracy over time.

This research project aims to extend the concept of data assimilation to urban CASoS. While complex systems models for CASoS have been developed, effectively integrating real-time data into these models – akin to the processes employed in weather forecasting – continues to present a significant challenge (Monti *et al.*, 2023; Asher *et al.*, 2023). Applying the principles of data assimilation to urban CASoS models, is a critical step toward developing accurate and responsive urban digital twins, ensuring that these sophisticated models reflect the dynamic nature of urban systems.

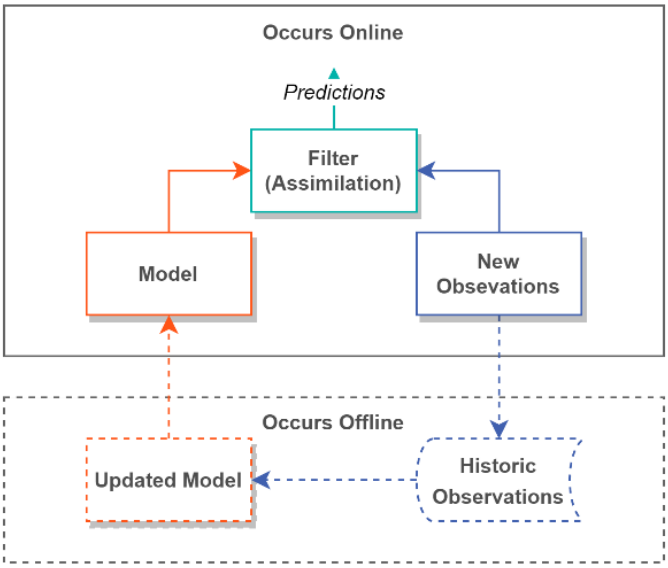

Figure 1 Data Assimilation Framework

1.2 Connected Digital Twins

Two essential characteristic of digital twins is that they can keep ultra-high synchronicity and fidelity with the physical space through metrology (measurements)—this process is shown in Figure 2 (Van der Valk *et al.*, 2020). Data assimilation is one method which can be used to achieve this. Synchronicity here implies at least a unidirectional real-time data link between the physical system to the model, and fidelity implies a model that contains sufficient information to be considered a faithful replica of the physical system (Jones *et al.*, 2020; Van der Valk *et al.*, 2020). These characteristics have demonstrated immense value in product lifecycle management (PLM) (Lim *et al.*, 2020) and increasingly, in building information modelling (BIM) (Boje *et al.*, 2020).

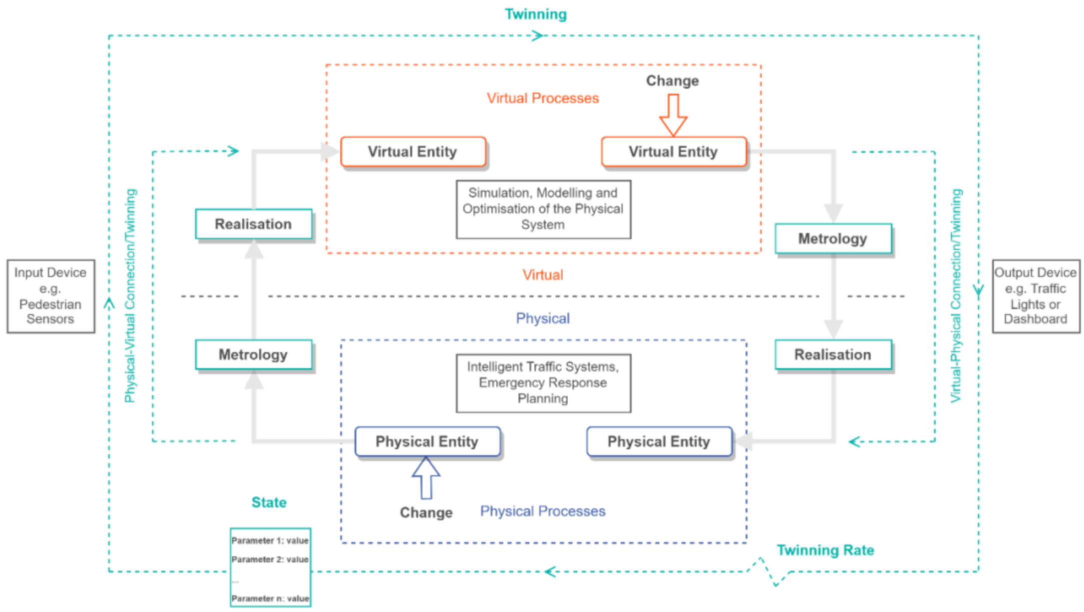

Figure** 2 **Physical/digital twinning cycle - based on Jones et al. (2020)

However, in both PLM and BIM, creating a digital twin typically involves models of *complicated* systems at the asset level—for example, the HVAC, power, water, structural and data systems for an office building modelled and optimised in real-time would be classified as a digital twin of a *complicated* system of systems (Baldwin *et al.*, 2011). However, a *complex* system is defined as one where understanding the overall behaviour necessitates knowledge of the individual actors and their interactions within the system (Baldwin *et al.*, 2011).

The integration of existing digital twins poses significant challenges, particularly when interfacing asset-level digital twins with synchronous models that represent complex human behaviours (Ivanov *et al.*, 2020; Digital Twin Hub, 2022; CDBB, 2022a). Initiatives such as the UK's Data and Analytics Facility for National Infrastructure (DAFNI), the Australian Urban Research Infrastructure Network (AURIN), and the UK Centre for Digital Built Britain (CDBB) are crucial in providing frameworks to ensure that new digital twins are complementary to existing ones through standardisation of common practices such as data management (CDBB, 2020; DAFNI, 2023). Creating and implementing data standards is currently a critical roadblock in creating a connected system of digital twins (Digital Twin Hub, 2022). The complexity of building *connected digital twins*, considering the variety of stakeholders, goals, and the intricate nature of cities, is exponentially harder than that of BIM and PLM (Digital Twin Hub, 2022; Lei *et al.*, 2023; Ivanov *et al.*, 2020).

The development of a city-scale connected digital twin requires a decentralised and evolutionary approach, underpinned by collaborative efforts, since no single entity can manage such a vast undertaking independently (Diakite *et al.*, 2022). Following this decentralised approach to development, the technological infrastructure of a city-scale digital twin can be split into two primary functions: digital twins for situational awareness (DTfSA), and digital twins for strategic decision-making (DTfSDM). The Gemini Principles (developed by the CDBB) separate these functions as follows (Bolton *et al.*, 2018):

- DTfSA - **"**A dynamic model of an asset, with input of current performance data from the physical twin via live data flows from sensors which feedback into the physical twin via real-time control".
- DTfSDM – "A static strategic planning model of a system, with input of long-term condition data from the physical twin via corporate systems which feedback into the physical twin via the capital investment process".

Critically, while these functions can be separated, a DTfSA can mature into a DTfSDM through re-analysis of the streamed data at a later date. The DTfSA during its lifetime can provide crucial partly simulated, partly observed data for the entire urban system through its coupling of real-time data input with complex modelling output. This allows strategic decision-making twins to utilise the significant volume of detailed historic data collected by DTfSA's.

The key difference between the function of a DTfSA and a DTfSDM lies in the use of real-time data. A DTfSA necessitates an accurate, current representation of the system, along with short-term predictions about its immediate future state, underscoring the critical role of real-time data. Barr *et al.* (2020a) for example, highlights how real-time data facilitates adaptive decision-making, while Shi *et al.* (2021a) discuss the impact of real-time data on predictive accuracy in complex systems. The city (considered as a CASoS) can be broken down into many systems (CAS) that exhibit unique behaviours and characteristics. Examples include economic and ecological systems, the spread of diseases, land-use planning, housing markets, voting behaviours, supply chain management, population dynamics, cultural evolution, the impact of climate change, public transportation, traffic flow, and healthcare systems. These are all systems renowned for their predictive challenges due to their inherent complexity. DTfSA can significantly enhance city planning and management services through greater understanding of behaviours in much the same way that the weather forecast allows for preparation against unexpected weather events (Hu *et al.*, 2021; Wang *et al.*, 2022). DTfSA could enable:

- Immediate response to anomalies (for instance, rapidly addressing traffic accidents and public transport delays).
- Dynamic adaptation (such as adaptive traffic control systems dynamically adjusting signal timings to optimise flow and minimise congestion).
- Reduced downtime and losses (for instance, promptly identifying and addressing issues in infrastructure, such as faults in metro lines).
- Optimisation of resources (for example, intelligent lighting systems that modulate brightness according to real-time ambient light levels and the presence of pedestrians or vehicles, enhancing energy efficiency and resource utilisation).
- Improved decision-making (for example, during natural disasters or accidents, real-time data providing detailed information on the location and severity of incidents enables emergency responders to swiftly and effectively deploy resources where needed).
- Customer satisfaction (for instance, providing public transport users with real-time updates on bus and train arrivals, enhancing the commuting experience).

The applications for real-time digital twins of complex systems are many and varied (CDBB, 2022b; Diakite *et al.*, 2022). One of the major challenges is in understanding what types of models can/should be coupled with real-time data streams to facilitate the predictive element of *situational awareness*. Situational awareness, defined by Endsley (1995) as the perception, understanding, and future projection of environmental elements, is crucial in dynamic decision-making. Systems characterised by decentralised interactions, diversity among entities, and emergent behaviour (like the systems mentioned above) not only fall into the CAS category but are also exceptionally well-suited for analysis through agent-based models (ABMs) (Epstein, 1999; Macal and North, 2009; Heard *et al.*, 2015; Railsback and Grimm, 2019).

1.3 Situational Awareness Digital Twin for Complex Systems

ABMs excel where multiple factors interact in intricate ways, leading to outcomes that are not straightforward to predict — often involving multi-scale feedback loops, non-linear relationships, and dependencies that can lead to emergent behaviours (Epstein, 1999). For example, traffic patterns in a city are influenced by road layouts, traffic signals, driver behaviour, public transportation systems, weather, school holidays, and a multitude of other factors. These factors lead to emergent behaviour such as traffic congestion; and affect travel times and accident rates in complex ways (Batty, 2013).

Whilst ABMs are incredibly effective at modelling these complex interactions, they do so offline, and the output of the ABM is typically very difficult to validate (Heppenstall et al., 2021). Challenges in real-time prediction and computational overhead in ABMs necessitate the development of efficient surrogate models (also referred to as *emulators* or *meta models*), that provide a computationally less intensive alternative to ABMs (Heppenstall and Malleson, 2020; Heppenstall *et al.*, 2021). While ABMs offer detailed insights into urban dynamics through their 'bottom-up' simulation approach, they face limitations in real-time processing and computational efficiency (Zhang, Li and Zhang, 2020). The technical challenge of creating a situational awareness digital twin in urban environments bears resemblance to the difficulty of weather forecasting using a sparse network of sensors (Zheng *et al.*, 2014; Batty, 2018). Mobility networks for example, are influenced by a range of human and environmental factors, from the weather to public holidays, to the latent effects of the covid pandemic on behaviour patterns. The variety of different factors makes the task of attaining high-resolution, short-term information and prediction for decision-making particularly difficult (Xu *et al.*, 2023).

While ABMs excel at simulating intricate interactions, integrating these models with real-time data remains largely uncharted (Epstein, 1999; Gilbert, 2019; Heppenstall and Malleson, 2020). Fundamentally, ABMs cannot take a series of inputs and make a prediction — their analytical power comes from creating agents – either by explicitly programming their behaviour or for more advanced models using machine learning to generate agent behaviours (Brearcliffe and Crooks, 2021) – and running the model to see how complex behaviour emerges as a result of the agents' interactions with their environment. Under existing modelling paradigms and available computational frameworks, a direct coupling between real-time and ABMs is not operationally feasible. However, taking inputs and creating predictions, is where machine learning methods excel. Innovative approaches like the application of machine learning to real-time data for creating ABM surrogates are emerging as a result (Kieu *et al.*, 2022). These surrogate models, simplified representations of more complex ABMs, can use data generated from a limited number of ABM simulations to predict system states, thus offering a solution to the validation challenges of ABMs, and the need for timely near real-time predictions for improved situational awareness (Heard et al., 2015).

An example is Melbourne's Activity and Agent Based Model which offers a comprehensive model of the cities mobility system (Infrastructure Victoria, 2021). Outputs of the ABM can be extracted at certain points, such as the number of pedestrians crossing a particular street. The ABM is used to generate a comprehensive dataset under a number of different initial conditions, through a number of simulations. The results of the simulation are extracted at certain points of the model, yielding time-series at these locations. The surrogate is then trained on this set of spatially distributed time-series data and learns the spatiotemporal patterns in the data. The trained model is then fed real data, for example data extracted from CCTV footage about the actual number of pedestrians at a location and makes a prediction about the number pedestrians at another nearby sensor. As the surrogate has been trained on the outputs of the ABM, the spatiotemporal dependencies should be captured in the surrogate, and the surrogate should be able to rapidly produce a prediction based on ABM states that are most similar to the real data input. The potential of surrogates in simulating complex urban systems in this way is significant, yet their application in large models using real data is still in its infancy (Malleson *et al.*, 2019; Ten Broeke *et al.*, 2021; Tang and Malleson, 2022; Ternes *et al.*, 2022).

Many deep-learning architectures, for example, Graph Neural Networks (GNNs) — which are well-suited for handling the inherent spatiotemporal data dependencies of urban dynamics — have shown promise in recent research into surrogates (Gilmer *et al.*, 2017; Jiang and Luo, 2022). For the first time to knowledge, this study aims to bring real-time data from a comprehensive urban sensor network together with appropriate modelling approaches (for example a hybrid deep-learning ABM surrogate) to make reliable, consistent, and scalable near real-time predictions of situational awareness.

2 Aim and Research Questions

2.1 Aim

Aim: To develop an AI system for the prediction of spatiotemporal dynamics in the built environment using near real-time geospatial IoT sensor data.

It is critical that computationally efficient models are developed to address the need for effective monitoring of complex systems. Enabling this technology unlocks significant economic and social value through faster identification of anomalous events (e.g. overcrowding of public spaces), more efficient use of resources (e.g. energy usage), and improved decision making (e.g. emergency response). Developing DTfSA will help cities to become more resilient to the impacts of climate change such as flooding by enabling faster response and intervention; safer through monitoring of threats such as overcrowding and spread of disease; more inclusive through the creation of open platforms that allow citizens to better understand the problems that exist in cities; and sustainable through more efficient resource allocation and eventually through enabling more effective strategic decision-making for meeting commitments such as net zero.

Three research questions have been identified to achieve this aim.

2.2 Research Questions

RQ1: *What is the utility and suitability of IoT sensors to capture the complex spatial and temporal dynamics of the urban environment in near real-time?*

Understanding the data that is available through exploratory analysis is a critical first step in developing DTfSA. Utility and suitability here refer to the effectiveness and appropriateness of the data respectively – both conditions need to be met to enable DTfSA. Effective data allows for a model to function correctly, and appropriate data allow for the outputs of a model to be explained.

RQ2: *Can machine-learning methods be used to develop predictive models of complex urban dynamics using near real-time sensor data?*

Measuring the effectiveness of machine-learning methods for predicting urban dynamics using historic sensor data is necessary to establish whether the data sources are effectively capturing the spatiotemporal patterns that exist in complex systems. This is critical - if the spatial distribution and temporal acquisition (data coverage) is inadequate then the surrogate is likely to make inaccurate predictions.

RQ3: *How can ABMs and the predictive outputs from research question 2 be combined within a data assimilation framework for improved understanding of urban dynamics?*

Enhancing the predictive capabilities of machine-learning models using the rich outputs of ABMs is the final goal of this research. This is based on the success of data-assimilation frameworks in creating real-time models of complex systems in other fields like meteorology, and the burgeoning field of research into creating surrogates of agent-based models.

3 Objectives and Methods

The four research objectives presented in this section, in addition to comprising the main chapters in the PhD thesis, are each intended to form a scientific publication.

3.1 *Objective 1*

To assess the quality of near real-time sensor data.

3.1.1 Overview

Central to the construction of predictive models is the comprehension of the data's quality — its suitability and utility — in predicting specific behaviours. This objective aims to devise methodologies for evaluating the quality of IoT data by employing various quality metrics. The research will build on the results of the MRes project and will entail the adaptation of an established data quality (DQ) taxonomy, applying it to IoT data streams from select pilot projects.

The relevance of both data utility and suitability varies according to the application domain. Data utility encompasses the overall effectiveness of the data, encapsulating attributes such as accuracy, completeness, and timeliness (Liu *et al.*, 2020; Fizza *et al.*, 2022; Mansouri *et al.*, 2023). Conversely, data suitability pertains to the data's compatibility with a specific purpose. This includes considerations like format and structural alignment with particular analytical tools or methodologies (what level of data preprocessing is necessary), as well as the data's contextual appropriateness, such as its congruence with the spatial and temporal aspects of the study (is the data from a similar enough time period, location, and culture to where it is being applied).

The investigation will encompass a series of case studies to evaluate the performance of the IoT data in these varied scenarios. For instance, one case study might involve the utilisation of pedestrian sensors to gauge overcrowding in a railway station. The research will explore multiple facets of DQ to ascertain the feasibility of making confident predictions about whether overcrowding thresholds are being exceeded. Key considerations include factors such as the station's maximum capacity and the reliability of the sensor data.

Another case study may focus on examining the suitability of data for augmenting existing anomaly detection methods in pedestrian or traffic contexts. This aspect of the research will critically assess the contextual relevance of the data, particularly its efficacy in developing models capable of detecting anomalies in diverse locations and times. Such an approach underscores the importance of not only the data's intrinsic qualities but also its adaptability to various situational demands, thereby offering valuable insights into its broader applicability.

3.1.2 Proposed Methods

Existing frameworks, such as those developed by Liu *et al.* (2020) or Mansouri *et al.* (2023), provide a robust structure for systematically evaluating each DQ aspect in the context of real-time data streams. Liu *et al.* (2020) identify definitions for six DQ dimensions based on existing literature. These are accuracy, timeliness, completeness, data volume, utility, and concordance. The authors also note the importance of future research into developing guidelines or checklists for DQ in IoT systems. Mansouri *et al.* (2023) identify twenty-five DQ dimensions, and ten metrics to measure DQ issues (such as redundancy, uncertainty, and ambiguity), that could serve as a checklist for DQ from IoT data sources.

A series of application case studies, such as the use of pedestrian sensors in railway stations, and anomaly detection in pedestrian or vehicle data, will serve as practical contexts for applying these dimensions and metrics to assess DQ. Data collection will be conducted using a combination of automated data extraction tools (for computing metrics such as completeness) and manual assessments (for properties of the sensors such as object detection accuracy calculated by the manufacturers).

DQ metrics, such as those developed by Fizza *et al.* (2022), will be utilised to measure the confidence in data based on its timeliness, particularly in applications requiring high situational awareness. Fizza *et al.* (2022) propose a model to compute age of IoT data that can be used by IoT application to cope with uncertainty in the data during the decision-making/actuation process. Data completeness will be measured in a similar fashion, following methods such as those presented by Ehrlinger and Wöß (2022) who calculate weighted or unweighted metrics for measuring the breadth, depth, and scope of information contained in the data completeness analogously to the accuracy metric on different aggregation levels with a weighted arithmetic mean. Confidence in accuracy — the degree to which observations of objects truly reflect their real-world situation — will need to be collected manually through engagement with manufacturers and operators of the IoT sensors.

The effect of DQ metrics like accuracy, timeliness, and completeness on predictions will be assessed through comparative model evaluation, using metrics such as root mean squared error (RMSE) and R squared. Popular models for making timeseries predictions, include statistical and numerical methods such as seasonal autoregressive integrated moving average (SARIMA) and optimised dynamic mode decomposition (ODMD); and neural network approaches such as transformers and long short-term memory units (LSTMs) will be investigated (Kutz *et al.*, 2016; Du *et al.*, 2020; Peppa *et al.*, 2021; Kieu *et al.*, 2022). A combination of these models will be used in this investigation following a detailed review of literature.

3.1.3 Deliverables

- Greater understanding of the DQ issues present in real-time IoT data.
- A new software library for the continuous monitoring of sensor data quality.

3.2 *Objective 2*

Assess the spatiotemporal dependency of near real-time sensor data.

3.2.1 Overview

This objective focuses on examining the spatiotemporal dependencies within IoT data streams from a network of sensors — this is critical for understanding whether sensor coverage is adequate which in turn depends on the DQ measures investigated in the previous objective. High spatiotemporal dependency would suggest adequate coverage, whereas an absence of dependency suggests additional systems are influencing behaviour that are not being measured, indicating additional sensors or additional data sources may be needed. This research will explore artificial intelligence methods for analysing spatially and temporally variable data, particularly in areas where continuous spatiotemporal fields are measured at irregular points in space, using sparse sensor networks like pedestrian or traffic CCTV cameras. By extending this investigation to non-Euclidian spatiotemporal fields such as pedestrian or vehicle networks, the study aims to deepen the understanding of spatial and temporal relationships in complex urban systems.

3.2.2 Proposed Methods

The methodology first involves an exploratory analysis to understand the expected lag in data patterns, such as estimating the walking, cycling, or driving time between sensors. The next step will be to train a preliminary deep learning model (e.g., a graph neural network) and assess its ability to make spatiotemporal predictions. The model will use sensor data that has been split into training/testing/validation subsets and evaluated using standard performance metrics.

An example case study might be a network of streets equipped with pedestrian sensors in the city centre. The role of exploratory analysis will be in analysing the similarity in data counts and the movement patterns they imply. The primary challenge anticipated is the limited spatial distribution, reliability, and availability of sensors. This initial research will be followed by an investigation of literature around preprocessing, to generate encodings that may help to improve the model performance, for example decomposing temporal signals from each IoT source using techniques like principal component analysis, Fourier analysis (Amato *et al.*, 2020), or optimised dynamic model decomposition are both approaches that could be investigated for this purpose (Kutz *et al.*, 2016).

To assess the similarity of spatial and temporal components, the study will use statistical methods such as spatial autocorrelation (Amato *et al.*, 2020) and dynamic time-warping (Froese *et al.*, 2020), alongside deep-learning approaches such as hybrid graph neural networks. Given the non-Euclidean nature of pedestrian networks, conventional deep learning approaches such as convolutional neural networks (CNNs) may not be suitable (Klemmer *et al.*, 2019). Instead, graph-neural networks and generative adversarial networks (GANs) with local autocorrelation, as indicated by recent research (Klemmer *et al.*, 2019; Jiang and Luo, 2022), will be explored.

3.2.3 Deliverables

- Greater understanding around the spatiotemporal dependency of IoT sensors.

- A series of notebooks containing the modelling and analysis results for a limited case study.
- A library of helper functions for manipulating spatiotemporal IoT data to use in future objectives.

3.3 *Objective 3*

Assimilate outputs of agent-based models with real-time sensor data in order to monitor urban systems in real-time.

3.3.1 Overview

This objective seeks to develop a surrogate model for simulating complex interactions in urban environments. The aim is to enable real-time processing of complex data through assimilation of ABM outputs with real-time data from IoT sensors. The surrogate model is essentially a computationally efficient stand-in (often referred to as a *surrogate*, *emulator*, or *meta*-model) for the ABM that enables rapid processing of data. This builds on existing work in this field (Lamperti *et al.*, 2018; Kieu *et al.*, 2022; Zhang *et al.*, 2020) where surrogate methods are tested with success on synthetic data. This objective will involve stakeholder collaboration with industry or government bodies to pre-emptively ensure the research adds practical value, and to facilitate the transition from theory to real-world application.

Statistical and machine learning models, although effective in predicting aggregate outcomes like footfall counts, often fall short in delivering deeper insights such as pedestrian density or delays that ABMs can provide (Kieu *et al.*, 2022). Moreover, the data-intensive nature of these simpler models pose a challenge, as they require extensive datasets to achieve versatility (Monti *et al.*, 2023). In many cases, certain system states might remain unobserved, limiting the ability of these models to make accurate predictions for scenarios not represented in the training data. This highlights the need for a more advanced approach that combines the comprehensive analytical power of ABMs with the efficiency of machine learning (Heppenstall *et al.*, 2021).

A machine-learning framework will be established that can effectively mimic the behaviour of ABMs by learning from their spatiotemporally rich outputs. These surrogates, effectively functioning as a computational bridge or assimilation filter, will translate the comprehensive data from ABMs into actionable insights with reduced processing time. The focus will be on creating a mapping between real-time aggregated IoT data and the outputs of ABMs, thus enabling more efficient and accurate simulations of urban dynamics (Kieu *et al.*, 2022).

3.3.2 Proposed Methods

The methodology involves a review of existing work in this area to identify potential enhancements of existing methods. The primary approach will be to train deep-learning models on the outputs of ABMs. The focus will be on developing models that can capture high-dimensional, nonlinear relationships inherent in urban dynamics. Various deep learning architectures will be explored for their suitability in approximating complex ABM outputs. This will follow on from the successful models identified in objective 2, like hybrid GNN-LSTM (Xu Zhang and Xia, 2022) and GANs with local autocorrelation (Amato *et al.*, 2020).

The performance of these surrogate models will be evaluated based on two components: firstly, their ability to replicate the behaviour observed in the ABM; and secondly, in their ability to make 'real-time' predictions of urban dynamics using historic IoT data as input. Figure 3 shows the overview of this process. The ABM is run offline to produce to a spatiotemporally rich training set (from a particular mobility type e.g., pedestrians), that is used to train the surrogate model. Sensors of a corresponding type to the training data produced by the ABM are then used to feed the deep learning surrogate with data upon which to make predictions. Continuing with the example of pedestrian sensors, the historic set of pedestrian data is used to evaluate the ability of the trained deep learning surrogate to make spatiotemporal predictions about pedestrians counts.

The network of sensors will be subset into testing and validation sets. The surrogate will be fed the test subset to make spatiotemporal predictions about the system, for different time periods *t.* The accuracy and reliability of the surrogate will then be validated using metrics like RMSE, MAE, and R squared. If the surrogate can accurately predict the value at these sensors for given time *t,* then spatiotemporal predictions can be tested in real-time applications and model evaluation will begin. Evaluation will involve measuring predictive performance, computational efficiency, robustness against different types of data (data from other sensor types - for example cyclists or motor vehicles if only pedestrians have been used up until this point), and its generalisability (for example IoT sensor networks from other cities). The time-taken to make predictions is critical, the model will need to be able to make a prediction in a timeframe that is useful, which will fundamentally limit the size of the surrogate used.

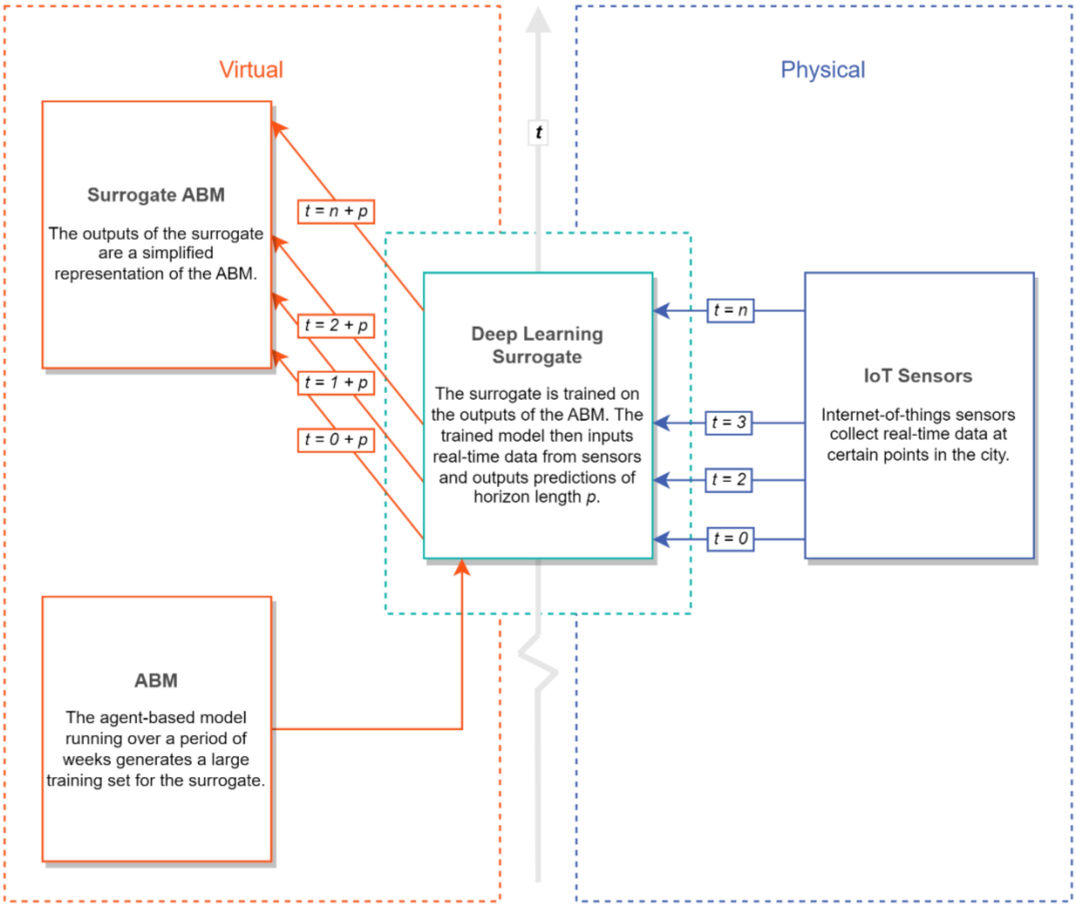

Figure 3 Methodological Overview - Research Objective 3

3.3.3 Deliverables

- Greater understanding around using agent-based models to train surrogate models.
- A surrogate model capable of real-time prediction with a demonstration for a particular case study.
- A pilot module of code for implementing real-time spatiotemporal predictions on IoT sensors.

3.4 *Objective 4*

Evaluate the approach using real-world case-studies and develop a roadmap for scalable deployment.

3.4.1 Overview

The real-world case studies used in this objective will result from engagement with academic partners. Likely candidates include working with Newcastle City Council alongside a national data science institute such as DAFNI or the Alan Turing Institute (ATI). The aim would be to use this technology to develop a critical piece of infrastructure such as emergency response capability. This research will include a 1-3 month placement in a city that has invested in IoT technology for recording urban dynamics, such as Singapore or Melbourne. The placement location would be used as another case study. Engagement with stakeholders will occur in the first month of this research project to identify the most critical use cases for this technology. This will include organising a 6-month internship at a partner institution such as DAFNI, ATI, or DSTL which could further inform the case studies investigated in this objective.

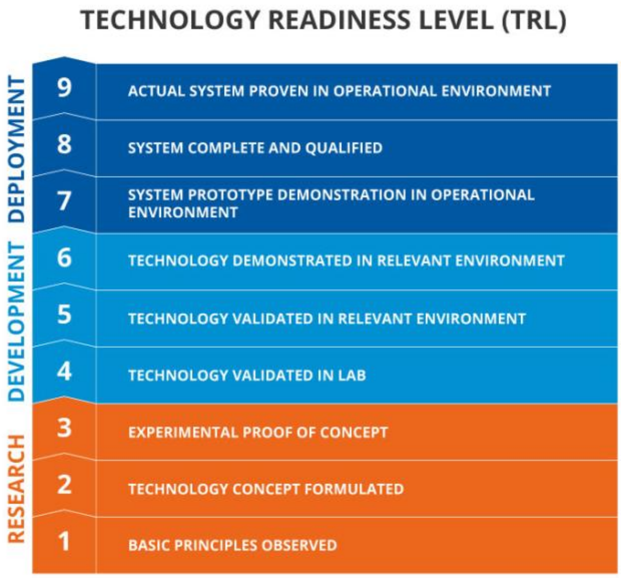

Figure 4 Technology Readiness Level

Technology readiness levels (TRL) shown in Figure 4 are a type of measurement system used to assess the maturity level of a particular technology (Manning, 2023). The final objective of this research seeks to look at how the outputs of previous objectives can be taken from pilot projects (at TRL 2 or 3) focusing on limited case-studies to a deployable piece of software that is fit for purpose (at TRL 5 or 6).

This might involve working with councils for monitoring overcrowding during large events, or with transport agencies for signalling maintenance requirements or automating traffic light timings. This objective will focus on issues found during model evaluation, for example scalability—how effectively the pilot models function when increasing the spatial and temporal complexity of the modelled system—and generalisability of the research—the application of the research to other cities, or urban environments, with different levels of IoT sensor deployment.

Critically, this research involves further engagement with stakeholders, building on the relationships developed during objective 3. Collaboration with other researchers will allow the methods to be tested on different ABMs to understand the generalisability and scalability of the objective 3 surrogate. This objective will begin with a reflective process to understand the limitations in wider deployment through evaluating the method on other ABMs and subsequently developing a roadmap for delivering DTfSA. Institutions like the Alan Turing Institute, Digital Twin Hub, and DAFNI are leaders in this space would be most beneficial to engage with. This objective offers an opportunity to coauthor with other researchers that are developing technology in this space to raise awareness of the work carried out.

3.4.2 Proposed Methods

Demonstrating scalability and generalisability will involve using the surrogate architecture on similar ABMs. Do the methods continue to work when increasing the number of sensors? At what spatial/temporal resolutions do the methods breakdown? Does the model generalise to other case studies? These questions will be used to imagine how the research might look like as an infrastructure service and offer a guide for any future work in this area.

Up until this point the surrogate will only have been tested on a single large ABM and a sensor network that corresponds to the ABMs purpose. Testing on similar ABMs, perhaps one that serves the same purpose but operates a different level of granularity, or is located in another city, is expected to highlight issues relating to generalisability and scalability. It is likely that accuracy will need to be sacrificed for computational efficiency for a more granular ABM that serves the same geographic area. Understanding where these trade-offs occur is critical for building a roadmap to DTfSA deployment. Developing this roadmap will involve the creation of a translation framework (identification of challenges such as ensuring the technology meets relevant industry standards and regulatory requirements; and identifying and mitigating risks associated with technology development etc.). The roadmap will include demonstrating how the surrogate integrates with existing digital infrastructure frameworks like DAFNI. It could also involve investigating methods for improving the computational run-time of the code, for example a rewrite into a low-level programming language like C, as well as engaging with stakeholders such as local councils to reflect on the utility of the prototype.

- An understanding of the limitations of real-time prediction in terms of utility and suitability of the developed methods for different data sources.
- A pilot software library for deploying situational awareness capabilities to a network of spatiotemporally dependent urban IoT sensors.
- A roadmap to higher technology readiness levels for the pilot software.

4 Novelty of the Proposed Research

This research seeks to develop new methods in making near real-time predictions for complex urban systems such as movements within pedestrian or vehicle networks derived from a network of non-independent sensors. Whilst there is a large amount of research into using data to understanding urban mobility, there is very little that looks at the challenges of creating decision making platforms that utilise real-time data. There is very little existing literature that focusses on developing methods for measuring dependency between IoT sensors. This work intends to build on the work of centralised urban data repositories (like the Urban Observatories in the UK), by investigating how a network of sensors might be used to make real-time predictions about the states of complex urban systems. This is in order to develop our understanding of complex behaviour and causality in complex systems allowing us to more effectively define and test solutions for complex and wicked problems. The research aims to contribute the following to the field of complex systems modelling:

- Greater understanding about the real dynamics of complex urban systems.
- An enhanced understanding of causality in complex urban systems.

And to deliver the following technical capabilities:

- A demonstration of value for the data collected by centralised open urban repositories.
- Pathways to making this data ‘AI ready’.
- A demonstration of a real-time cloud-based web-app that provides useful information derived from the sensor data that can be used for improved decision making.
- A package of code that conforms to best-practice and is built to be compatible with digital urban-twin frameworks such as DAFNI/Gemini.
- A series of publications showcasing any scientific advancements made by this research.

5 References

Alizadeh, R., Allen, J. K. and Mistree, F. (2020) 'Managing computational complexity using surrogate models: a critical review', *Research in Engineering Design,* 31, pp. 275-298.

Amato, F., Guignard, F., Robert, S. and Kanevski, M. (2020) 'A novel framework for spatio-temporal prediction of environmental data using deep learning', *Scientific reports,* 10(1), pp. 22243.

Asher, M., Oswald, Y. and Malleson, N. (2023) 'Predicting Pedestrian Counts using Machine Learning', *AGILE: GIScience Series,* 4, pp. 18.

Baldwin, W. C., Felder, W. N. and Sauser, B. J. (2011) 'Taxonomy of increasingly complex systems', *International Journal of Industrial and Systems Engineering,* 9(3), pp. 298-316.

Barr, S., Johnson, S., Ming, X., Peppa, M., Dong, N., Wen, Z., Robson, C., Smith, L., James, P. and Wilkinson, D. (2020a) 'Flood-prepared: A nowcasting system for real-time impact adaption to surface water flooding in cities', *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences,* 6, pp. 9-15.

Barr, S. L., Johnson, S., Ming, X., Peppa, M., Dong, N., Wen, Z., Robson, C., Smith, L., James, P., Wilkinson, D., Heaps, S., Laing, Q., Xiao, W., Dawson, R. and Ranjan, R. (2020b) 'Flood-Prepared: A Nowcasting System for Real-Time Impact Adaption to Surface Water Flooding in Cities', *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences,* VI-4/W2-2020, pp. 9-15.

Basadur, M., Conklin, J. and VanPatter, G. (2007) 'Rethinking wicked problems', *Conversation with G vanPatter) NextD Journal,* 10.

Batty, M. (2018) 'Digital twins', 45(5), pp. 817-820.

Batty, M. (2022) 'The post-pandemic city: speculation through simulation', *Cities,* 124, pp. 103594.

Birks, D., Heppenstall, A. and Malleson, N. (2020) 'Towards the development of societal twins', *ECAI 2020*: IOS Press, pp. 2883-2884.

Boje, C., Guerriero, A., Kubicki, S. and Rezgui, Y. (2020) 'Towards a semantic Construction Digital Twin: Directions for future research', *Automation in construction,* 114, pp. 103179.

Bolton, A., Butler, L., Dabson, I., Enzer, M., Evans, M., Fenemore, T., Harradence, F., Keaney, E., Kemp, A. and Luck, A. (2018) 'Gemini principles'.

Brearcliffe, D. K. and Crooks, A. 'Creating intelligent agents: Combining agent-based modeling with machine learning'. *Proceedings of the 2020 Conference of The Computational Social Science Society of the Americas*: Springer, 31-58.

Caparros-Midwood, D., Dawson, R. and Barr, S. (2019) 'Low carbon, low risk, low density: Resolving choices about sustainable development in cities', *Cities,* 89, pp. 252-267.

CDBB (2020) *The approach to delivering a National Digital Twin for the United Kingdom*: Centre for Digital Built Britain.

CDBB (2022a) *Gemini Papers: How to Enable an Ecosystem of Connected Digital Twins*: Centre for Digital Built Britain.

CDBB (2022b) *Gemini Papers: Why connected digital twins?* : Centre for Digital Built Britain.

DAFNI (2023) *Data & Analytics Facility for National Infrastructure to advance UK infrastructure research*. Available at: https://www.dafni.ac.uk/about-2/ (Accessed: 27/10/2023.

Dawson, R. J. (2011) 'Potential pitfalls on the transition to more sustainable cities and how they might be avoided', *Carbon Management,* 2(2), pp. 175-188.

Diakite, A., Ng, L., Barton, J., Rigby, M., Williams, K., Barr, S. and Zlatanova, S. (2022) 'Liveable City Digital Twin: a Pilot Project for the City of Liverpool (nsw, Australia)', *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences,* 10, pp. 45-52.

Digital Twin Hub (2022) *Digital Twin Roadblocks - Final Report*: Digital Twin Hub.

Du, B., Peng, H., Wang, S., Bhuiyan, M. Z. A., Wang, L., Gong, Q., Liu, L. and Li, J. (2020) 'Deep Irregular Convolutional Residual LSTM for Urban Traffic Passenger Flows Prediction', *IEEE Transactions on Intelligent Transportation Systems,* 21(3), pp. 972-985.

Ehrlinger, L. and Wöß, W. (2022) 'A survey of data quality measurement and monitoring tools', *Frontiers in big data,* 5, pp. 850611.

Endsley, M. R. (1995) 'Toward a Theory of Situation Awareness in Dynamic Systems', *Human Factors,* 37(1), pp. 32-64.
Epstein, J. M. (1999) 'Agent-based computational models and generative social science', *Complexity,* 4(5), pp. 41-60.
Fizza, K., Jayaraman, P. P., Banerjee, A., Georgakopoulos, D. and Ranjan, R. 'Age of data aware Internet of Things applications'. *2022 IEEE 19th Annual Consumer Communications & Networking Conference (CCNC)*: IEEE, 399-404.
Froese, V., Jain, B., Niedermeier, R. and Renken, M. (2020) 'Comparing temporal graphs using dynamic time warping', *Social Network Analysis and Mining,* 10, pp. 1-16.
Gilbert, N. (2019) *Agent-based models.* Sage Publications.
Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O. and Dahl, G. E. 'Neural message passing for quantum chemistry'. *International conference on machine learning*: PMLR, 1263-1272.
Goodspeed, R. (2015) 'Smart cities: moving beyond urban cybernetics to tackle wicked problems', *Cambridge Journal of Regions, Economy and Society,* 8(1), pp. 79-92.
Grieves, M. (2022) 'Intelligent digital twins and the development and management of complex systems', *Digital Twin,* 2(8), pp. 8.
Heard, D., Dent, G., Schifeling, T. and Banks, D. (2015) 'Agent-based models and microsimulation', *Annual Review of Statistics and Its Application,* 2, pp. 259-272.
Heppenstall, A., Crooks, A., Malleson, N., Manley, E., Ge, J. and Batty, M. (2021) 'Future Developments in Geographical Agent-Based Models: Challenges and Opportunities', *Geogr Anal,* 53(1), pp. 76-91.
Heppenstall, A. and Malleson, N. 'Building cities from slime mould, agents and quantum field theory'. *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020)*: Leeds.
Hu, C., Fan, W., Zeng, E., Hang, Z., Wang, F., Qi, L. and Bhuiyan, M. Z. A. (2021) 'Digital twin-assisted real-time traffic data prediction method for 5G-enabled internet of vehicles', *IEEE Transactions on Industrial Informatics,* 18(4), pp. 2811-2819.
Infrastructure Victoria (2021) *Transporting Melbourne's Recovery*.
Ivanov, S., Nikolskaya, K., Radchenko, G., Sokolinsky, L. and Zymbler, M. 'Digital twin of city: Concept overview'. *2020 Global Smart Industry Conference (GloSIC)*: IEEE, 178-186.
Jiang, W. and Luo, J. (2022) 'Graph neural network for traffic forecasting: A survey', *Expert Systems with Applications,* 207, pp. 117921.
Jones, D., Snider, C., Nassehi, A., Yon, J. and Hicks, B. (2020) 'Characterising the Digital Twin: A systematic literature review', *CIRP journal of manufacturing science and technology,* 29, pp. 36-52.
Kieu, M., Nguyen, H., Ward, J. A. and Malleson, N. (2022) 'Towards real-time predictions using emulators of agent-based models', *Journal of Simulation*, pp. 1-18.
Klemmer, K., Koshiyama, A. and Flennerhag, S. (2019) 'Augmenting correlation structures in spatial data using deep generative models', *arXiv preprint arXiv:1905.09796*.
Knapp, R. (2015) 'Wholesome Design for Wicked Problems'. Available at: https://www.publicsphereproject.org/content/wholesome-design-wicked-problems.
Kutz, J. N., Brunton, S. L., Brunton, B. W. and Proctor, J. L. (2016) *Dynamic mode decomposition: data-driven modeling of complex systems.* SIAM.
Lamperti, F., Roventini, A. and Sani, A. (2018) 'Agent-based model calibration using machine learning surrogates', *Journal of Economic Dynamics and Control,* 90, pp. 366-389.
Lei, B., Janssen, P., Stoter, J. and Biljecki, F. (2023) 'Challenges of urban digital twins: A systematic review and a Delphi expert survey', *Automation in Construction,* 147, pp. 104716.
Lewis, J. M., Lakshmivarahan, S. and Dhall, S. (2006) *Dynamic data assimilation: a least squares approach.* Cambridge University Press.
Lim, K. Y. H., Zheng, P. and Chen, C.-H. (2020) 'A state-of-the-art survey of Digital Twin: techniques, engineering product lifecycle management and business innovation perspectives', *Journal of Intelligent Manufacturing,* 31, pp. 1313-1337.

Liu, C., Nitschke, P., Williams, S. P. and Zowghi, D. (2020) 'Data quality and the Internet of Things', *Computing,* 102(2), pp. 573-599.
Macal, C. M. and North, M. J. 'Agent-based modeling and simulation'. *Proceedings of the 2009 winter simulation conference (WSC)*: IEEE, 86-98.
Malleson, N., Minors, K., Kieu, L.-M., Ward, J. A., West, A. A. and Heppenstall, A. (2019) 'Simulating crowds in real time with agent-based modelling and a particle filter', *arXiv preprint arXiv:1909.09397*.
Manning, C. G. (2023) 'Technology Readiness Levels', *Space Communications Navigation Program*, Available at: NASA. Available at: https://www.nasa.gov/directorates/somd/space-communications-navigation-program/technology-readiness-levels/#:~:text=Technology%20Readiness%20Levels%20(TRL)%20are,based%20on%20the%20projects%20progress.
Mansouri, T., Sadeghi Moghadam, M. R., Monshizadeh, F. and Zareravasan, A. (2023) 'IoT data quality issues and potential solutions: a literature review', *The Computer Journal,* 66(3), pp. 615-625.
Monti, C., Pangallo, M., De Francisci Morales, G. and Bonchi, F. (2023) 'On learning agent-based models from data', *Scientific Reports,* 13(1), pp. 9268.
Peppa, M. V., Komar, T., Xiao, W., James, P., Robson, C., Xing, J. and Barr, S. (2021) 'Towards an End-to-End Framework of CCTV-Based Urban Traffic Volume Detection and Prediction', *Sensors (Basel),* 21(2), pp. 629.
Railsback, S. F. and Grimm, V. (2019) *Agent-based and individual-based modeling: a practical introduction.* Princeton university press.
Rittel, H. W. and Webber, M. M. (1973) 'Dilemmas in a general theory of planning', *Policy sciences,* 4(2), pp. 155-169.
Shi, W., Goodchild, M. F., Batty, M., Kwan, M.-P. and Zhang, A. (2021a) *Urban Infomatics.* Springer.
Shi, W., Goodchild, M. F., Batty, M., Kwan, M.-P. and Zhang, A. (2021b) *Urban informatics.* Springer.
Tang, D. and Malleson, N. (2022) 'Data assimilation with agent-based models using Markov chain sampling', *arXiv pre-print server*.
Ten Broeke, G., van Voorn, G., Ligtenberg, A. and Molenaar, J. (2021) 'The use of surrogate models to analyse agent-based models', *Journal of Artificial Societies and Social Simulation,* 24(2).
Ternes, P., Ward, J. A., Heppenstall, A., Kumar, V., Kieu, L.-M. and Malleson, N. (2022) 'Data assimilation and agent-based modelling: towards the incorporation of categorical agent parameters', *Open Research Europe,* 1(131).
Tonkinwise, C. (2015) 'Design for Transitions–from and to what?', *Design Philosophy Papers,* 13(1), pp. 85-92.
UN-Habitat (2022) *World Cities Report 2022*: United Nations Habitat.
United Nations (2015) *Goal 11: Make cities inclusive, safe, resilient and sustainable*: United Nations.
Van der Valk, H., Haße, H., Möller, F., Arbter, M., Henning, J.-L. and Otto, B. 'A Taxonomy of Digital Twins'. *AMCIS*.
Wang, Z., Gupta, R., Han, K., Wang, H., Ganlath, A., Ammar, N. and Tiwari, P. (2022) 'Mobility digital twin: Concept, architecture, case study, and future challenges', *IEEE Internet of Things Journal,* 9(18), pp. 17452-17467.
Xu, H., Berres, A., Yoginath, S. B., Sorensen, H., Nugent, P. J., Severino, J., Tennille, S. A., Moore, A., Jones, W. and Sanyal, J. (2023) 'Smart Mobility in the Cloud: Enabling Real-Time Situational Awareness and Cyber-Physical Control Through a Digital Twin for Traffic', *IEEE Transactions on Intelligent Transportation Systems,* 24(3), pp. 3145-3156.
Xu Zhang, S. W., Liang Yan, Jiangfan Feng, and and Xia, Y. (2022) 'A Hybrid-Convolution Spatial–Temporal Recurrent Network For Traffic Flow Prediction', *The Computer Journal,* 00(00).
Zhang, Y., Li, Z. and Zhang, Y. (2020) 'Validation and calibration of an agent-based model: A surrogate approach', *Discrete Dynamics in Nature and Society,* 2020.
Zheng, Y., Capra, L., Wolfson, O. and Yang, H. (2014) 'Urban computing: concepts, methodologies, and applications', *ACM Transactions on Intelligent Systems and Technology (TIST),* 5(3), pp. 1-55.

6 Appendices